

Delay Performance in IP Routers

Patrik Carlsson, Doru Constantinescu, Adrian Popescu, Markus Fiedler and Arne A. Nilsson

Dept. of Telecommunication Systems

School of Engineering

Blekinge Institute of Technology

371 79 Karlskrona, Sweden

{patrik.carlsson, doru.constantinescu, adrian.popescu, markus.fiedler, arne.nilsson}@bth.se

Abstract

The main goals of the paper are towards an understanding of the delay process in best-effort Internet for both non-congested and congested networks. A dedicated measurement system is reported for delay measurements in IP routers, which follows specifications of the IETF RFC 2679. The system is using both passive measurements and active probing. Dedicated application-layer software is used to generate UDP traffic with TCP-like characteristics. Pareto traffic models are used to generate self-similar traffic in the link. The reported results are in the form of several important statistics regarding processing delay of a router, router delay for a single data flow, router delay for more data flows as well as end-to-end delay for a chain of routers. We confirm results reported earlier about the fact that the delay in IP routers is generally influenced by traffic characteristics, link conditions and, at some extent, details in hardware implementation and different IOS releases. The delay in IP routers usually shows heavy-tailed characteristics. It may also occasionally show extreme values, which are due to improper functioning of the routers.

Keywords: traffic measurements, delay performance, IP routers, traffic self-similarity

1 Introduction

As the Internet emerges as the backbone of worldwide business and commercial activities, end-to-end (e2e) Quality of Service (QoS) for data transfer becomes a significant factor. End-to-end delay is a key metric in evaluating the performance of networks as well as the quality of service perceived by end users. Today network capacities are being deliberately overengineered in the Internet so that the packet loss rate is very low. Throughput maximization can be done by minimizing the e2e delay. However, given the heterogeneity of the network and the fact that the overengineering solution is not adopted everywhere, especially not by backbone teleoperators in developing countries, the question arises as to how the delay performance impacts the e2e performance. There are several important parameters that may impact the e2e delay performance in the link, e.g., traffic self-similarity, routing flaps and link utilization [9, 11].

Several papers report on e2e delay performance, and both Round-Trip Time (RTT) and One-Way Transit Time (OWTT) are considered [2, 5, 11, 12]. Traffic measurements based on both passive measurements and/or active probing are used. Generally, RTT measurements are simpler but the analysis is more complex due to different problems related to clock synchronization, packet timestamping, protocol complexity, asymmetries in direct and return paths as well as path variations [5]. Other problems related to difficulties in measuring queuing delays in operational routers and switches further complicates the picture [11].

As a general comment, it has been observed that both RTT and OWTT show large "peak-to-peak" variations, in the sense that maximum delays far exceeds minimum delays. Further, it has been observed that OWTT variations (for opposite directions) are asymmetric in most cases, with different delay distributions, and they seem to be correlated with packet loss rates [12]. Periodic delay spikes and packet losses have been also observed, which seem to be a consequence of routing flaps [11]. Typical distributions for OWTT have been observed to have a Gamma-like shape and to possess heavy tail [2, 10]. The parameters of the Gamma distribution have been observed to depend upon the path (e.g., regional, backbone) and the time of the day.

The main goals of the paper are towards an understanding of the delay process in best-effort Internet for both non-congested and congested networks. We have designed a measurement system to do delay measurements in IP routers, which follows specifications of the IETF RFC 2679 [1]. The system uses both passive measurements and active probing. Dedicated application-layer software is used to generate UDP traffic with TCP-like characteristics. The well-known interactions between TCP sources and network are thus avoided. UDP is not aware of any network congestion, and this gives us the choice of doing experiments where the focus is on the network only. The software consists of a client and a server running on two different hosts, which are separated by a number of routers. Pareto traffic models are used to generate self-similar traffic in the link. Both packet inter-arrival times and packet sizes are matching real traffic models. A passive measurement system is used for data collection that is based on using several so-called Measurement Points (MPs), each of them equipped with DAG monitoring cards [4, 7]. The combination of passive measurements and active probing, together with using the DAG monitoring system, gives us an unique possibility to perform precise traffic measurements as well as the flexibility needed to compensate for the lack of analytic solutions.

The rest of the paper is organized as follows. In Section 2 we provide a short review of One-Way Transit Time together with the associated delay components. In Section 3 we describe the measurement system and the technology used to collect data. We discuss the accuracy and the limitations as well. Section 4 is devoted to the set of experiments done, and we report on specific details related to these experiments. Section 5 is dedicated to reporting the results obtained on delay performance. Finally, Section 6 concludes the paper.

2 Delay Components

One-Way Transit Time (OWTT) is measured by timestamping a specific packet at the sender, sending the packet into the network, and comparing then the timestamp with the timestamp generated at the receiver [1]. Packet timestamping can be done either by software (for the case of delay measurements at the application level) or by hardware (for the case of delay measurements in the network), and in this case special hardware is used. Clock synchronization between the sender and the receiver nodes is important for the precision of one-way delay measurements. On top of that, delay measurements at the application level are sensitive to possible uncertainties related to the difference between the "wire time" and the "host time" as well.

OWTT has several components:

$$OWTT = D_{prop} + \sum_{i=0}^N D_{n,i} \quad (1)$$

where the delay per node i , $D_{n,i}$ is given by:

$$D_{n,i} = D_{tr,i} + D_{proc,i} + D_{q,i} \quad (2)$$

The components are as follows:

- D_{prop} is the total propagation delay along the physical links that make up the Internet path between the sender and the receiver. This time is solely determined by the properties of the communication channel as well as the distance. It is independent of traffic conditions on the links.
- N is the number of nodes between the sender and the receiver.
- $D_{tr,i}$ is the transmission time for node i . This is the time it takes for the node i to copy the packet into the first buffer as well as to serialize the packet over the communication link. It depends on the packet length and it is inversely proportional to the link speed.
- $D_{proc,i}$ is the processing delay at the node i . This is the time needed to process an incoming packet (e.g., to decode the packet header, to check for bit errors, to lookup routes in a routing table, to recompute the checksum of the IP header) as well as the time needed to prepare the packet for further transmission, on another link. This delay depends on parameters like network protocol, computational power at node, and efficiency of network interface cards.
- $D_{q,i}$ is the queueing delay in the node i . This delay refers to the waiting times in buffers, and depends upon traffic characteristics, link conditions (e.g., link utilization, interference with other IP packets) as well as implementation details of the node.

Several statistics, e.g., mean, median, maximum, minimum, variance, peakedness and probability distribution function, are usually used in the description of delay for non-corrupted packets. Typical values obtained for OWTT range from tens of μs (between two hosts on the same LAN) to hundreds of ms (in the case of hosts placed in different continents) [3].

For a general discussion, the OWTT delay can be partitioned into two components, a deterministic delay D_d and a stochastic delay D_s :

$$OWTT = D_d + D_s \quad (3)$$

D_{prop} , D_{tr} and (partly) D_{proc} are contributing to the deterministic delay D_d , whereas the stochastic delay D_s is created by D_q and, at some extent, D_{proc} . The stochastic part of the router processing delay can be observed especially in the case of low and very low link utilization, when the queueing delays are minor. There are several parameters that affect the stochastic delay D_s , and the most important are the link utilization (ρ), traffic characteristics (as described by the Hurst parameter H) and node parameters (including the well-known routing flaps).

3 Delay Measurements

We are reporting delay measurements done at the network level [6]. Figure 1 shows the measurement configuration used in our experiments. The key component in the system is a Measurement Point (MP), the device that does the actual packet capturing. The capabilities of an MP are decided by the capture hardware that is installed in the MP, and in our experiments we use the DAG 3.5E network monitoring card [7]. The MPs are capable of collecting and timestamping frames with an accuracy of less than 100 ns. Data analysis is done off-line. Furthermore, the MPs are locally synchronized to each other.

The networks that we are measuring are 10 Mbps full duplex Ethernets. On a 10 Mbps Ethernet the maximum frame rate is 14881 frames/s, and this equals a frame interarrival time of 67.2 μs . This time is significantly larger than the timestamp accuracy of the MP.

The routers R1, R2 and R3 are all of the same type (Cisco 3620). The source host A, the sink host E and the hosts that generate cross traffic B, C, and D are all identical with regards to hardware and software configuration.

The MPs were instructed to collect the first 96 bytes of every frame received on the Ethernet links. This way we are able to collect link, network and transport headers as well as part of the payload. Frames that are smaller than 96 bytes are zero padded.

To estimate the delay that a packet experiences we need to first identify the packets as they pass the MPs on their way through the routers. Hashing is used for the identification and matching of packets. The hashing function is implemented with the SHA-1 Secure Hash Algorithm. All captured packets are masked before hashing. The hash covers the entire IP header including the source and destination IP addresses, the IP header identification field and other fields, except the Time To Live (TTL) and the Header Checksum fields as they are changed at every router. 37 bytes of the IP payload (including IP options and eventual padding) are included in the hash as well. The identification and matching software provides timestamp readings for every two adjacent MPs. Further details on hashing and packet identification as well as analysis of duplicate and unmatched packets are reported in [6].

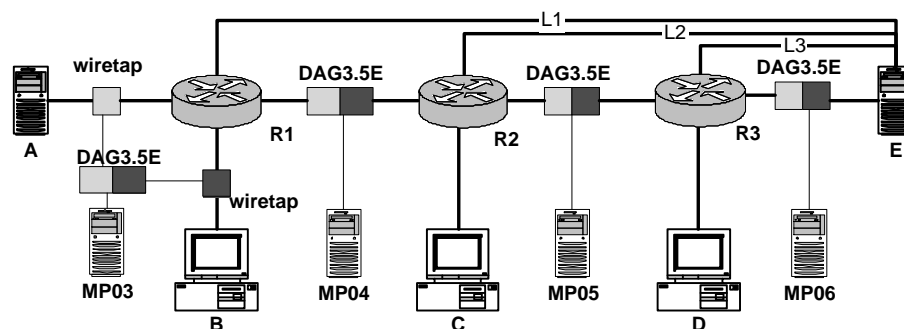


Figure 1: Measurement configuration

4 Experiments

A dedicated application-layer software is used to generate UDP traffic with TCP-like characteristics [6]. The traffic generated between the source A and the sink E follows a Pareto distribution for the packet length with the shape parameter α (which determines the mean and the variance) and the location parameter β (which determines the minimum value). An exponential distribution with parameter λ is used for inter-packet gap times and the (measured) link utilization ρ depends on λ . This model matches well traffic models observed for World Wide Web, which is one of the most important contributors to the traffic in Internet [8]. Higher traffic intensities (than those measured in real networks) are considered in our experiments as well, with the consequence of higher loads on routers, and thus we obtain better delay models.

The cross traffic generated by the traffic generators B, C, D has a form that approaches fractional Brownian motion (fBm) traffic with self-similar characteristics, which is typical for Ethernet [9]. This traffic is generated in a similar way, i.e., Pareto distribution for packet sizes and exponential distribution for inter-packet times. The difference however is that a large number of processes are used in this case to generate a large number of Pareto traffic flows in every traffic generator. This gives us the choice of doing experiments where we can control diverse parameters of the traffic mixture in the link and especially the Hurst parameter and the link utilization ρ .

Table 1 shows the summary of experiments and the associated parameters. In experiment 1, only the source computer A generates traffic, with diverse characteristics. One process is used to generate Pareto traffic with α parameter shown in table and a traffic intensity λ that corresponds to the (measured) ρ values shown in table. Further, the parameter $\beta = 40$. Nine traces have been generated in this case, which are shown as **1-1**, **1-2**, to **1-9**.

Table 1: Summary of experiments and traffic generation parameters

Exp ^a	A α, ρ (1)	Exp ^b	A α, ρ (1)	B α, ρ (100)	Exp ^c	A α, ρ (1)	B α, ρ (50+50)	C α, ρ (50+50)	D α, ρ (50+50)
1-1	2, 0.2	2-1	2, 0.2	1.2, 0.1	3-1	2, 0.2	1.2, 0.1	1.2, 0.1	1.2, 0.1
1-2	2, 0.4	2-2	2, 0.4	1.2, 0.1	3-2	2, 0.4	1.2, 0.1	1.2, 0.1	1.2, 0.1
1-3	2, 0.6	2-3	2, 0.6	1.2, 0.1	3-3	2, 0.6	1.2, 0.1	1.2, 0.1	1.2, 0.1
1-4	1.6, 0.2	2-4	1.6, 0.2	1.2, 0.1	3-4	1.6, 0.2	1.2, 0.1	1.2, 0.1	1.2, 0.1
1-5	1.6, 0.4	2-5	1.6, 0.4	1.2, 0.1	3-5	1.6, 0.4	1.2, 0.1	1.2, 0.1	1.2, 0.1
1-6	1.6, 0.6	2-6	1.6, 0.6	1.2, 0.1	3-6	1.6, 0.6	1.2, 0.1	1.2, 0.1	1.2, 0.1
1-7	1.2, 0.2	2-7	1.2, 0.2	1.2, 0.1	3-7	1.2, 0.2	1.2, 0.1	1.2, 0.1	1.2, 0.1
1-8	1.2, 0.4	2-8	1.2, 0.4	1.2, 0.1	3-8	1.2, 0.4	1.2, 0.1	1.2, 0.1	1.2, 0.1
1-9	1.2, 0.6	2-9	1.2, 0.6	1.2, 0.1	3-9	1.2, 0.6	1.2, 0.1	1.2, 0.1	1.2, 0.1

^aExperiment 1^bExperiment 2^cExperiment 3

In experiment 2, both computers A and B are generating traffic. The traffic generated by computer A has the same characteristics as the traffic generated in experiment 1. Computer B generates fBm-like traffic with $H \approx 0.9$ and $\rho \approx 0.1$. One hundred processes are used for traffic generation in computer B. Nine traces have been generated in experiment 2, which are shown as **2-1**, **2-2**, to **2-9**.

For experiment 3, the computer A is still generating traffic with the same characteristics as in experiment 1. The difference is that now all three computers B, C and D are generating fBm-like traffic flows, with $H \approx 0.9$ and $\rho \approx 0.1$. One hundred processes are used for traffic generation in every computer. The generated traffic flows are broken up in the routers R1, R2 and R3 in a way that 50 % of every flow is merging the traffic coming from computer A and the rest of 50 % (for every flow) is crossing the routers only. Nine traces have been generated in experiment 3 as well, which are shown as **3-1**, **3-2**, to **3-9**. Figure 2 shows examples of traces collected at the computers A, B, C, D, and E, together with the associated histograms.

5 Delay Performance

5.1 Processing Delay of a Router

Figure 3 shows typical processing delays (D_{proc}) in a router, which are measured for IP packets containing ICMP and UDP payloads, and for variable payload sizes (between 32 and 1450 bytes). The associated histograms are reported in the figures 4 (for ICMP payload) and 5 (for UDP payload), respectively. Cisco 3620 routers have been used for this experiment. A number of 10000 samples have been generated for every payload size. The ICMP tests were done by using *ping*, and no options were used that would have required extra processing burden for the router. Very low traffic intensities were used (in the order of 10 frames/s) such as to avoid the presence of queueing delay in the router.

The mean delay for UDP samples was found to be 97.9 μ s, with a minimum of 74.2 μ s and a maximum of 1861.2 μ s. The mean delay for ICMP samples was found to be higher (as expected), i.e., 101 μ s. The minimum delay for ICMP samples was found to be 76.5 μ s and the maximum 1958.8 μ s.

We have also done similar experiments on other routers (e.g., Cisco 1605, Cisco 2514). We observed that the processing delays look similar, the difference however is that the delays may variate with up to about 100 μ s compared to the values reported above. This is clearly a difference that depends upon details in hardware implementation and different IOS releases.




Figure 2: Traffic collected in the measurement set-up and the associated histograms

5.2 Router Delay for a Single Data Flow

In experiment 1 the first router R1 is receiving traffic from the computer A only. The Measurement Points MP03 and MP06 are used to capture the traffic. The routers R2 and R3 are not used and the MP04 and MP05 are not connected either. Nine traces have been generated, which correspond to different values for burstiness and traffic intensities of the generated traffic. Every generated trace is quite big and it has about one million packets. The reason for that is because of the well-known problems related to the slow convergence to steady-state in the case of heavy-tailed workloads [9].

Table 2 reports the main statistics of the router delay measured between the sink computer E and the generator A. The most obvious feature that we observe is related to the rather limited disparity of these results. The traces show that packets experience delays that have quite similar statistics, with mean and variance values that are slightly increasing with H and ρ . Further, the experiment is also showing that the number of samples with large delays is quite low. Most of samples have maximum




Figure 3: Example of router processing delay for ICMP and UDP payloads




Figure 4: Router processing delay for ICMP payloads with different sizes




Figure 5: Router processing delay for UDP payloads with different sizes

Table 2: Summary of OWTT results for experiment 1

Experiment	Mean	95% Conf.	Variance	Max	Min	Peakedness	Dups ^a / Unm ^b
1-1	0.205	± 0.0002335	0.01420	3.98	0.13	0.069	0 / 0
1-2	0.241	± 0.0002895	0.02182	1.79	0.12	0.091	0 / 0
1-3	0.287	± 0.0006054	0.09539	55.5	0.12	0.33	0 / 0
1-4	0.227	± 0.0003033	0.02395	2.08	0.12	0.11	0 / 0
1-5	0.267	± 0.0003604	0.03381	6.63	0.13	0.13	0 / 0
1-6	0.317	± 0.0004015	0.04195	1.97	0.13	0.13	0 / 0
1-7	0.262	± 0.0004293	0.04798	3	0.13	0.18	0 / 0
1-8	0.301	± 0.0004874	0.06184	3.62	0.13	0.21	0 / 0
1-9	0.349	± 0.0005405	0.07604	2.11	0.13	0.22	0 / 0

^aTotal number of duplicate packets

^bTotal number of unmatched packets

delays that are less than 7 ms, and we believe that they are due to queueing at the output link. There is however one delay (in the trace **1-3**), which has a maximum that is unusual large (55.5 ms) compared to the rest of traces. This is likely not created by queueing, but due to other reasons like the router stopping forwarding packets for a while. Further, the associated histograms (not reported in the paper) have been observed to have a Gamma-like shape with a heavy tail that is dependent on α and ρ .

5.3 Router Delay for More Data Flows

In experiment 2 the router R1 is receiving traffic from both computers A and B. In this case, the computer A is generating non-fBm-like traffic with the same characteristics like in experiment 1. The computer B is generating fBm-like traffic. The routers R2 and R3 are still not used and the MP04 and MP05 are not connected either.

As a result, nine traces have been generated, with different ρ for the traffic observed at the sink computer E. Every trace is about one million packets long. The Hurst parameter of the bit rate has been measured to be $H = 0.9$ for most of traces. There are however several traces showing larger values for H (e.g., $H = 1.09$) for large ρ and low α . These are traces corresponding to traffic with infinite mean. We are aware of this problem, we wanted however to test the router under very severe conditions. We therefore did not excluded these traces from our experiments.

Table 3 reports the main statistics of the router delay measured between the sink computer E and the generator A. We observe in this case that the delays show a larger disparity in the range of statistics. Though the mean values are quite similar, the variance and the peakedness show however a larger

Table 3: Summary of OWTT results for experiment 2

Experiment	Mean	95%-Conf.	Variance	Max	Min	Peakedness	Dups ^a / Unm ^b
2-1	0.758	± 0.003186	2.640	27.8	0.12	3.5	0 / 463
2-2	1.39	± 0.005123	6.808	34.9	0.12	4.9	0 / 2792
2-3	1.82	± 0.006031	9.421	33.8	0.13	5.2	0 / 3999
2-4	0.723	± 0.002944	2.253	36.4	0.13	3.1	0 / 211
2-5	0.855	± 0.003588	3.347	33.7	0.13	3.9	0 / 607
2-6	1.7	± 0.006046	9.467	44.4	0.13	5.6	0 / 3887
2-7	0.873	± 0.003843	3.822	49.3	0.12	4.4	0 / 4912
2-8	1.41	± 0.005840	8.854	51.4	0.13	6.3	0 / 1861
2-9	3.62	± 0.009987	25.68	55.8	0.13	7.1	0 / 9997

^aTotal number of duplicate packets

^bTotal number of unmatched packets

disparity. Furthermore, this experiment shows that the number of samples with large delays is pretty high, and we believe that most of them are due to queueing. No delay is observed that is extremely large compared to the rest. The associated histograms have been observed to have a Gamma-like shape with a heavy tail that is dependent on α and ρ . Figure 6 shows an example of delay trace and the associated histogram obtained in experiment **2-1**.

5.4 End-to-End Delay for a Chain of Routers

In experiment 3 the entire setup shown in figure 1 is used. Computer A generates the same traffic pattern as in the previous experiments. Computers B, C and D are generating both merging traffic (50 %) and crossing traffic (50 %). The crossing traffic enters on the same router port as the merging traffic, but leaves on a port different from the port used by the A to E traffic stream.

Table 4 reports the main statistics of the router delay measured between the sink computer E and the generator A. The main observation in this case is regarding the large disparity for all statistics collected, except for the minimum. We also observe a large number of samples with large delays, and we believe that most of them are due to queueing in routers. We do not observe any sample with extremely large delay, the conclusion of which is that all three routers seem to work properly. The associated histograms have been observed to have a Gamma-like shape with a heavy tail. The tail has been observed to look similar for most of traces and it is dependent on α and ρ . Figure 7 shows an example of delay trace together with the associated histogram obtained in the experiment **3-7**.

Finally, the figure 8 shows the delay performance (mean and variance) obtained for all three experiments. It is observed the strong dependence of the mean and the variance on traffic characteristics and link utilization. Further, it is also observed that it is the α parameter of the generated traffic that mostly influences the router behavior at large link utilization. This is of course the consequence of queueing behavior, which typically shows heavy-tailed delay performance (e.g., Weibull distribution) in the case of traffic with Long-Range Dependence.

6 Conclusions

A measurement study of delays through IP routers has been reported. A dedicated passive measurement system has been designed to collect high-quality data traces and to measure e2e delay. The measurement set-up is reported in a detailed way together with the set of experiments done. Several important statistics have been reported about the delay for a router and for a chain of routers. Our results confirm other results reported earlier that the delay in IP routers is generally influenced by traf-

Table 4: Summary of OWTT results for experiment 3

Experiment	Mean	95%-Conf.	Variance	Max	Min	Peakedness	Dups ^a / Unm ^b
3-1	2.59	± 0.008419	18.32	70	0.40	7.1	0 / 7180
3-2	4.88	± 0.012140	37.36	92.2	0.41	7.6	0 / 25657
3-3	7.56	± 0.013240	43.56	77.7	0.41	5.8	0 / 44961
3-4	2.58	± 0.008468	18.56	77.3	0.41	7.2	0 / 5880
3-5	3.02	± 0.009749	24.53	80.9	0.41	8.1	0 / 8496
3-6	3.02	± 0.009749	24.53	80.9	0.41	8.1	0 / 51138
3-7	3.13	± 0.010160	26.69	102	0.41	8.5	0 / 6199
3-8	5.1	± 0.014930	56.96	107	0.41	11	0 / 18238
3-9	16	± 0.020870	100.2	110	0.42	6.3	0 / 115763

^aTotal number of duplicate packets

^bTotal number of unmatched packets




Figure 6: Delay trace and the associated histogram obtained in the experiment 2-1




Figure 7: Delay trace and the associated histogram obtained in the experiment 3-7

fic characteristics, link conditions and, at some extent, details in hardware implementation and IOS releases.

Our future work will be about modeling of the obtained results, to include possible correlations existent between queue lengths at adjacent nodes as well. Our goal is to find a formula for the end-to-end delay in a chain of routers. This will be finally validated in real networks, under conditions of real traffic.

References

- [1] Almes G., Kalidindi S. and Zekauskas M., *A One-way Delay Metric for IPPM*, IETF RFC 2679, 1999.
- [2] Bovy C. J., Mertodimedjo H. T., Hooghiemstra G., Uijterwaal H. and Van Mieghem P., *Analysis of End-to-End Delay Measurements in Internet*, ACM PAM, Fort Collins, Colorado, USA, 2002.
- [3] Caida 2001 Network Measurement Metrics WG., <http://www.caida.org/outreach/metricswg/faq.xml> [checked June 2004]
- [4] Carlsson P., Ekberg A. and Fiedler M., *On an Implementation of a Distributed Passive Measurement Infrastructure*, COST279 TD(03)042, 2003.

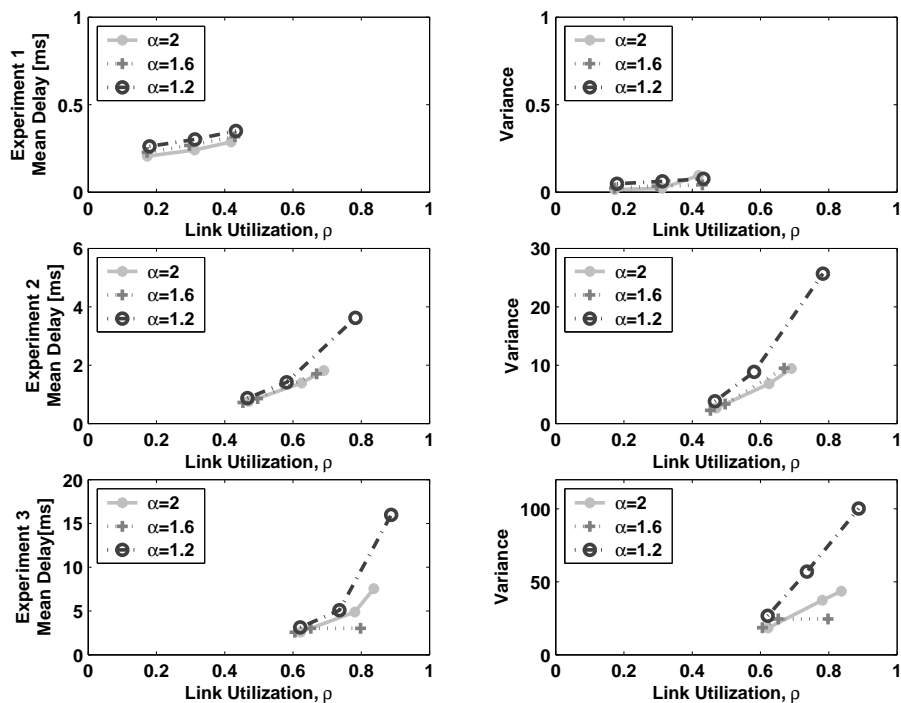


Figure 8: Summary of delay performance

- [5] Claffy K. C., Polyzos G. C. and Braun H. W., *Measurement Considerations for Assessing Unidirectional Latencies*, Journal of Internetworking, Vol. 4, No. 3, 1993.
- [6] Constantinescu D., Carlsson P., *One-Way Transit Time Measurements*, Technical Report, Blekinge Institute of Technology, 2004, ISSN:1103-1581.
- [7] *Endace Measurement Systems*, <http://www.endace.com>
- [8] Jena A. K., Popescu A. and Nilsson A. A., *Modeling and Evaluation of Internet Applications*, International Teletraffic Congress ITC-18, Berlin, Germany, 2003.
- [9] Leland W. E., Taqqu M. S., Willinger W. and Wilson D. V., *On the Self-Similar Nature of Ethernet Traffic (Extended Version)*, IEEE/ACM Transactions on Networking, Vol. 2, No. 1, 1994.
- [10] Mukherjee A., *On the Dynamics and Significance of Low Frequency Components of Internet Load*, Internetworking: Research and Experience, Vol. 5, 1994.
- [11] Papagiannaki K., Moon S., Fraleigh C., Thiran P. and Diot C., *Measurement and Analysis of Single-Hop Delay on an IP Backbone Network*, IEEE Journal on Selected Areas in Communications, Vol. 21, No. 5, August 2003.
- [12] Paxson V., *Measurements and Analysis of End-to-End Internet Dynamics*, PhD Dissertation, University of California at Berkeley, 1997.